Resource Discovery Systems at the UNT LIBRARIES
Phase Two Action Plan

UNT LIBRARIES

July 31, 2013

Authored by:
Mark Phillips – Assistant Dean for Digital Libraries
Neena Weng – Head, User Interfaces Unit
Jason Thomale – Resource Discovery Systems Librarian, User Interfaces Unit
William Hicks – Technical Coordinator, User Interfaces Unit

Executive Summary

Our 2011 report, *Resource Discovery Systems at the UNT Libraries*, provides a four-phase vision and model for how we should develop our resource discovery infrastructure to meet challenges our library faces in the 21st century. It concludes with a detailed plan for Phase One, which stipulates that we'd revisit the model and create a plan for Phase Two once we'd finished Phase One.

In the first phase, we redesigned two major aspects of our discovery infrastructure—the library catalog and the library website—both based on existing user data and feedback, and both with positive results. We acquired Serials Solutions' Summon product and implemented it as a *Find Articles* search, integrating it into our new website, which more than doubled usage of full-text articles over previous years. Finally, we implemented a new Integrated Library System (ILS)—Innovative Interfaces' Sierra. With the first phase now complete, this document outlines our Phase Two plan.

Upon revisiting our 2011 resource discovery system (RDS) implementation model, we have found that the second phase as presented needs revision. Our original model aimed, by the end of Phase Four, to build a unified, locally-developed discovery application atop an infrastructure of modular components that exchange data openly—and phases two through four would take us through a series of steps focused first on opening up our backend systems *before* building the application. Though our ultimate goal has not changed, our experiences with Phase One have taught us that this is not necessarily the best path. By taking advantage of techniques and lessons learned in Phase One (pages 3-4), we can instead focus on providing a better, more consistent user experience earlier in the plan without compromising our original vision and still work to improve the openness of our backend systems incrementally to prepare for future improvements. Deploying an application that powers a single-search tab on our library's search box and provides bento-box style search results (pages 5-8) will be Phase Two's hallmark.

Toward these ends, our Phase Two Action Plan recommends the following four objectives (beginning on page 9):

1. Continue making iterative updates to existing resource discovery and delivery services, interfaces, and systems to improve end-users' overall experience.
2. Conduct research and gather data needed to help us better understand user behavior, evaluate our systems, and further inform resource discovery system development.
3. Continue to work with other UNT Libraries departments to improve data quality.
4. Deploy an application to power a single-search tab on the library website that delivers a bento-box-style results display.

We plan to complete our Phase Two objectives by August 2015, after which we will revisit the model again and plan for Phase Three.

Contents

Changes to the RDS Implementation Model 3
Lessons Learned from Phase One 3
The Revised Model: Single Search and the "Bento Box" Display 5
Objectives 9
Objective 1: Continue making iterative updates to existing resource discovery and delivery services, interfaces, and systems to improve end-users' overall experience. 9
Objective 2: Conduct research and gather data needed to help us better understand user behavior, evaluate our systems, and further inform resource discovery system development. 10
Objective 3: Continue to work with other UNT Libraries departments to improve data quality. 12
Objective 4: Deploy an application to power a single-search tab on the library website that delivers a bento-box-style results display. 13
Timeline 14
References 16

Changes to the RDS Implementation Model

In 2011 we authored a comprehensive report about the then-current state of library resource discovery with the goal of providing a research- and data-based plan for improving how users discover our resources on the UNT Libraries website. The report outlines an overarching vision—our RDS Implementation Model—to help keep us on track. We divided the model into four separate phases, and we created a detailed action plan for Phase One and possible goals for the other three. At the end of our Phase One Action Plan we recommended that, because the RDS landscape is evolving quickly, we should reexamine the state of the RDS art after completing Phase One and then revise our vision and model as needed (Phillips, et al., 2011).

Now that we have finished Phase One, we have found that we were correct in making such a recommendation. Circumstances *have* evolved, and, as we have gained experience and user data from our Phase One efforts, we have gained a better understanding of how to go about accomplishing what we originally set out to accomplish.

In our original model, phases two through four imply a step-by-step, first-this-then-that plan, where the unified resource discovery interface depicted in Phase Four becomes possible once we've met the architectural goals outlined in phases two and three. Reevaluating following Phase One, however, has shown us that this division of work is actually too clear-cut and in some ways conflicts with our stated iterative design philosophy. We can, in fact, integrate and improve our discovery interfaces in degrees, implementing features that make the biggest difference to users and then refining them as we go based on usage data and user feedback, while simultaneously making architectural improvements to open up our data, improve system modularity, and enable us to tighten system integration in future iterations. In effect, this puts a looser version of the Phase Four unified resource discovery interface—manifested in the "single search box" for resource discovery—into our plan sooner than originally anticipated. Other architectural details that we once thought would be vital, such as a discovery-layer index for the catalog and integration of our Web-Scale Discovery service via an API, seem less crucial now for overall discovery service interface integration due to other strides we've made in Phase One. Although we can and should certainly experiment with such components, we no longer see them as our only option for accomplishing our goals.

Lessons Learned from Phase One

Affordances of Partially Open Architectural Components

Our original model presents architectural components as being either vendor-controlled (closed) or locally-controlled (open), but reality is not so unequivocal—and, although we state that some components may only be partially closed, the model itself does not really take this into account. As a result, it places unrealistic requirements on components being fully open source and/or locally-developed to meet our Phase Four goals.

For example, the model requires that we implement a discovery layer on top of our existing catalog in order to open up our catalog data and allow us to build new front-end interfaces for it. While an open

discovery layer or fully open-source catalog may be the ideal, it is not strictly necessary. Through our work in Phase One we have found that Sierra's more-open architecture affords us flexibility in how we query and extract data from the ILS that we didn't have with Millennium—specifically, Sierra's SQL interface allows us programmatic access to any records and transactional information stored in Sierra's database, through which our own software can interact directly with the database in real time. Although access to the SQL interface is read-only, and we thus cannot write data to Sierra this way, Innovative Interfaces, Inc. stated at the latest Innovative Users Group Conference that the company plans to release APIs beginning in Winter 2013 that will allow both read and write access, such a Resource Discovery and a Holds API (Jung & Jones, 2013). These possibilities make it more tenable to use Sierra directly to interact with locally-developed components without need for a fully open system. Although more openness would give us more flexibility in how we develop our local solutions, we should at least remain flexible and fully explore our options before committing to one particular system architecture.

Opening Up Closed Web-based Components

Related to the previous point, we have also learned that closed Web-based components are rarely actually completely closed. The Web is built on an open network, with many open technologies working in tandem to provide what end-users see. When necessary, locally-developed scripts and tools can intercept and parse Web pages generated by a vendor-controlled system to force otherwise closed components to behave more openly.

Much of our Phase One work on improving resource discovery interfaces involved customizing vendor-provided Web-based systems, where we had more or less flexibility to customize the look, feel, and functionality of the interface depending on the system. Whether the system is more restrictive (such as our e-Journal Portal) or less (such as the catalog interface), we found that we could employ custom client-side scripting to rewrite portions of the display in the browser before the browser renders the page. We used this workaround extensively and to great effect when we customized the library catalog display; we used it when we integrated our e-Journal Portal with our electronic resources interface in the catalog; and we used it when we added functionality to Summon to allow users to report broken full-text article links to us.

We have found that these and similar techniques can provide a viable, less resource intensive alternative to building something from scratch—at least as an incremental step forward if not a permanent solution. Considering such options opens up the possibility to provide better user-facing services more quickly while we continue working on developing our underlying system architecture.

Content Management System (CMS) as Discovery Platform

Until we redesigned the library website in Phase One we did not fully grasp the role that an open source library CMS could play in the overall discovery infrastructure, except as a system that powers the library website. The system we chose—Drupal—is in fact a fully featured, fully open source development platform that provides us additional options as we work to integrate the pieces of our discovery infrastructure with our library website. Drupal enjoys wide adoption amongst libraries, and the development of library-related Drupal modules means we could easily build off what others have

already done. For instance, an experimental Drupal module created by the University of Michigan that provides Summon integration could provide the basis for a bento-box style search results display (Bertram, 2011). (See: *The Revised Model: Single Search and the "Bento Box" Display*, below).

Resource Delivery

Allowing users to find resources that they can't obtain does them a disservice. Finding something listed in one of our discovery systems or on our website with nothing clearly indicating how to get it—or, worse, with a link to an item's full-text that doesn't actually resolve correctly—is incredibly frustrating. Our experience with Summon, the Web-Scale Discovery service that we implemented in Phase One, illustrates this perfectly. The majority of problems and negative user feedback we've had with Summon has been related to resource *delivery*, where users can search the system but then cannot obtain the items they find. Broken links to full-text articles are widespread enough that, during Phase One, we instituted a simple form allowing users to report them to us so that we can then have them reported to the vendor for resolution.

This also illustrates that roadblocks hampering delivery are not always under our control, and we thus are limited to some degree in how much we can improve them. In fact, we've identified at least three such limitations. First, as we learned from Summon, data provided and controlled by vendors can have issues that affect delivery that we can't directly resolve. Second, locally-controlled data can also have issues that are not within our direct sphere of influence or are otherwise difficult to fix: for instance, catalog records may contain URLs subject to link rot. Third, as with discovery services, delivery services may be tied up in third-party systems that are only partially customizable—such as the link resolver. Taking steps toward improving delivery will require that we attempt to mitigate or work around these limitations where possible.

The Revised Model: Single Search and the "Bento Box" Display

Phase Four of our original model describes a possible single-search solution. It features a unified, locally-developed application that sits on top of all of our resource discovery data stores—UNT-owned and otherwise—and communicates with each via APIs or other means. As already mentioned, it also relies on implementing a discovery layer index to interface with the library catalog as well as other local data. The original model is intentionally vague about the solution's interface. It implies one (and only one) search box that searches everything using asynchronous, federated-search-reminiscent techniques; it also implies that some (if not all) discovery activities can take place within the unified application rather than kicking users out to each native application.

Broadly speaking, the overarching vision behind the original model hasn't actually changed—a unified, locally-developed application is still one of our future goals. However, based on what we have learned in Phase One—as described in the previous section—we have clarified our ideas about how our single-search application should look and how it will function. We have learned that in Phase Two we can begin working to implement an initial version of this application that lets us provide users with a single-search option while still relying on customized versions of each native interface to accommodate the bulk of discovery activities (as in Figure 1, below).

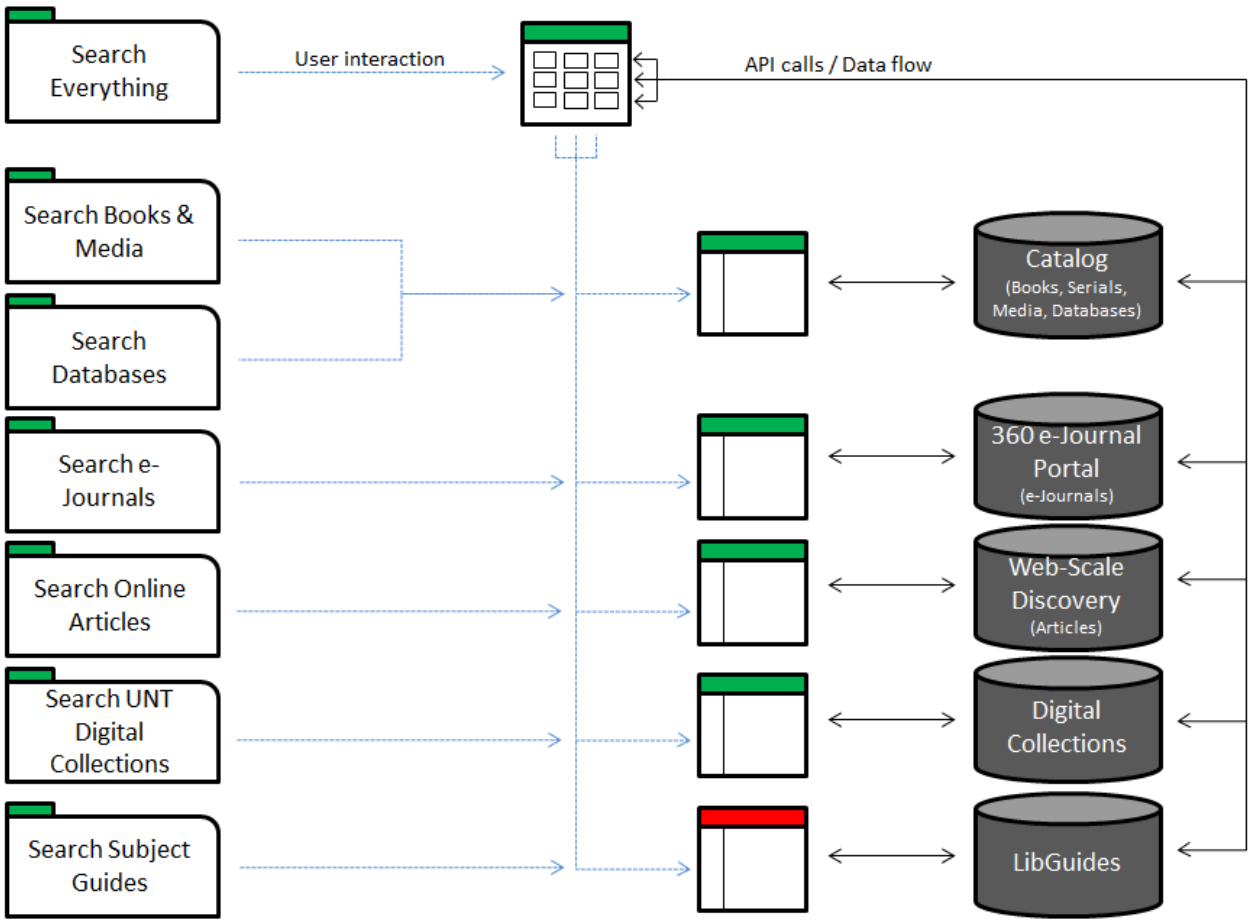

FIGURE 1: UNT LIBRARIES' UPDATED RDS VISION, PHASE TWO: Note, in this diagram, that the dotted blue lines represent user interaction and the solid black lines represent the flow of data in backend systems. The green, tabbed folders represent search tabs on our library website, the green and red boxes that resemble webpage wireframes represent applications, and the dark gray cylinders represent individual data stores within our infrastructure. Applications in green are those with interfaces that we have developed, will develop, or that we have or will have customized heavily; those in red are those that will still mostly use the stock vendor interface.

As Figure 1 shows, our revised Phase Two model adds a new tabbed search option whereby a user can search across multiple data stores with one query. The application we plan to build to power this new search will distribute the user's query to each included system, listen for results, and dynamically build a page that displays the top few results from each system in a separate pane on the page (with links that allow users to access each set of full results in its native system). For simplicity, we call this style of single-search implementation "bento box," after terminology that Tito Sierra coined while at North Carolina State University, as the multi-pane layout resembles a Japanese bento takeout box (Dempsey, 2012; Rochkind, 2012; Lown, Sierra, & Boyer, 2013). Figure 2, below, shows a prime example of a bento-box display.

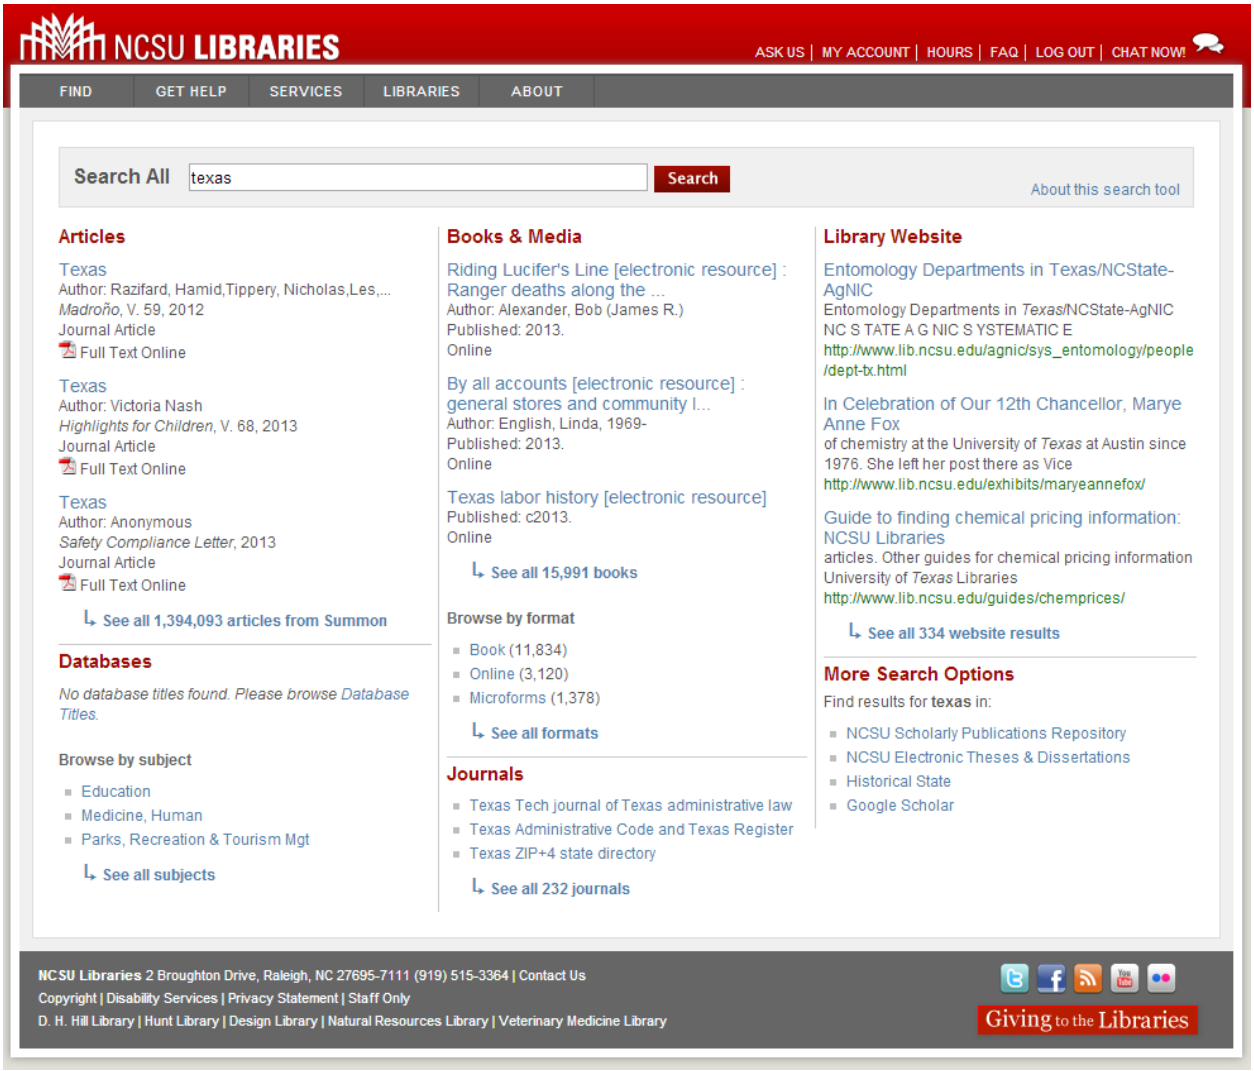

FIGURE 2: NCSU'S BENTO-BOX DISPLAY

Bento-box search displays are a departure from the more common Google-like, blended-results displays that most single-search library applications employ and have at least a couple of distinct advantages. First, generating a blended-results display requires all resource metadata to be contained within one index. Article-level metadata presents a problem, since article-level discovery is the exclusive realm of library vendors (as discussed in our RDS report), and vendors have a vested interest in ensuring that their data remains closed (Phillips, et al., 2011). This means that—to include article-level results in your single-search system—you *have to* import your library-controlled data (catalog records, digital library records, etc.) into the vendor's index and utilize the vendor's product to power your single-search application. Bento-box displays, on the other hand, function more like federated search—they query each index separately and thus allow us more easily to rope off vendor-indexed data while simultaneously allowing us to retain control over the application itself. Second, lumping all of the metadata for all content types together into one index and attempting to perform meaningful relevance ranking can be problematic when content types are as heterogeneous as those in most libraries. In such a system, the sheer volume of certain content types can drown out others. Plus, usability data is

ambivalent about whether users actually want article-level results to be combined with catalog results—some published user data casts doubt on the efficacy of combined results. Bento-box displays by design get around these problems because they give each type of content coming from each system its own pane in the display to ensure that nothing gets drowned out or lost. In a position paper prepared for use at Johns Hopkins University, Jonathan Rochkind (2012) explores in detail the benefits that a bento-box display provides over a blended-results display and comes to similar conclusions.

Take note that, in our revised model, the single-search box does *not* replace the other tabbed search options. Rather, we retain the individual tabbed searches that interact with particular systems—the catalog, our Serials Solutions e-Journal portal, Summon, the Digital Collections, and Lib Guides. User data that we collected during our website redesign project in Phase One convincingly demonstrated that users prefer to have options when searching library resources. Although they may not mind having the option to search all (or most) library resources at once, they also like having other options presented to them so that they can see what's available and easily search something more specific if needed (Weng, 2012).

Finally, as with the vision and implementation model presented in our original RDS report, we do not intend the revised model to shackle us as circumstances change. The diagram is an illustration, composed with an eye toward readability and visual symmetry. In reality, details about—for instance—exactly which components will be included in the bento box display will depend on an analysis to be done during Phase Two. Furthermore, in the first Phase Two objective, we explore iterative improvements we can make to backend systems that would let us retain the interface/application layer shown in the model while continuing to work toward open solutions such as what we originally envisioned. Such solutions may change certain details as they appear in the diagram while not substantially changing the overall idea that the diagram illustrates.

Objectives

Objective 1: Continue making iterative updates to existing resource discovery and delivery services, interfaces, and systems to improve end-users' overall experience.

Actions

- Upgrade our Summon instance to version 2.0 by January 30th, 2014.
- Improve the Summon interface to integrate the Find Online Articles service more seamlessly into the look, feel, and flow of the library website.
- Implement holds and resource requesting in the catalog to help users obtain physical library resources more easily by September 2013.
- Customize or redesign the Link Resolver interface to improve the overall user experience surrounding resource delivery services.
- Plan and test replacing the Google search appliance on the library website with Apache Solr.
- Continue investigating means to improve searching, browsing, display of, and end-user interaction with library catalog data.
- Continue working with III in our role as Sierra Early Adopter helping make recommendations for improvements to Sierra (and updating our system to incorporate improvements when available).

Summary

Phase One of our RDS Implementation Model comprised implementing several new systems to enhance our resource discovery environment: Summon to power a find-articles search, Sierra to help us begin improving the infrastructure surrounding the catalog, and Drupal to provide a more robust Web content management system. Now that these systems are part of our infrastructure, in Phase Two we can focus on iterative improvement, development, and integration of the interfaces and services that sit atop these systems.

We can also turn toward an aspect of the resource discovery experience that we did not touch upon in Phase One: resource *delivery*. Our trials with broken full-text article links in Summon have taught us that resource discovery is useless without equally robust resource delivery—e.g., helping users actually obtain the resources that they find, whether those resources are physical ones found via the catalog or virtual ones found via Summon.

The steps that we propose in this objective we hope will help us weld together our disparate resource discovery and delivery services to make them seem more as parts of one service than the amalgam that they still are today.

Benefits

- Updating to Summon version 2.0 will let us take advantage of the new features and enhancements that the Summon development team has made based on the rich usage and usability data it has collected over the years.
- Consistency is one of the hallmarks of usable interface design. Our plans will improve consistency in the following ways:
  - Redesigning the Summon interface so that the Find Online Articles service is better incorporated into the library website will help provide a more seamless experience for users of that service.
  - Redesigning the Link Resolver interface will help us make the resource delivery experience more seamless, and it will allow us the opportunity to consolidate delivery services, helping users obtain resources more easily.
  - Integrating holds and requesting into the catalog will allow us to further consolidate delivery services and make them function more consistently across our discovery services.
- Replacing our Google search appliance with Apache Solr would allow us more flexibility in customizing how our website search works and what is searched.
- A new Solr implementation could be used for more than just the website—we could adapt it to use for indexing catalog data, for example.
- Continuing to work with III as a Sierra Early Adopter affords us the opportunity to have feedback from our library incorporated into Sierra development. It also affords us the opportunity to gain early access to new Sierra features such as APIs.

Objective 2: Conduct research and gather data needed to help us better understand user behavior, evaluate our systems, and further inform resource discovery system development.

Actions

- Continue collecting and analyzing usage data across all library resource discovery systems, improving collection methods as needed to obtain more relevant data.
- Continue collecting and analyzing user feedback about library resource discovery systems.
- Provide technical support to the Collection Management Division for implementing the ARL MINES survey tool, which will help us better study our users who access electronic resources and gauge how they access them.
- Plan and implement studies exploring the pathways our users take to arrive at information resources to help us understand how we might inject our discovery and delivery services into their workflows to better serve them.
- Plan and implement studies exploring how users browse information and how we can take advantage of their browsing habits to supplement and enhance our search interfaces to improve resource discovery.

Summary

Every step of our RDS Implementation Plan has been informed by data. Conclusions in our original RDS report were supported by an extensive literature review plus user data collected from previous years' LibQual surveys and data collected from our peer institutions' RDS implementations. Our website redesign incorporated results from participatory design studies, user surveys, Web Content Workgroup feedback, and a comprehensive survey of other academic libraries' website components. At each step we have put into place mechanisms that gather usage data and end-user feedback (for the library catalog, for Summon, for the website, and for the electronic resources interface), and we have consulted this data as we've formulated what we want to accomplish in Phase Two.

Clearly we want to continue to collect as much data as we can about our users and how they use our systems to ensure that our actions continue to serve them, and we also want to improve our data collection mechanisms—such as our code that sends data to Google Analytics or Piwik—to ensure that we're collecting the data that most effectively answers our questions. Toward this end, we will support the Collection Management Division as they implement the MINES survey tool from the Association of Research Libraries. Data collected via this tool will help us obtain a better grasp on how users access our electronic resources. Beyond this, we have identified a couple of specific areas that we have so far neglected that we think would help us address particular shortcomings in future phases: finding out how we might improve discovery and delivery services for our patrons who may not necessarily visit the

library website, and finding out how we might incorporate interfaces that better support browsing to improve serendipitous resource discovery.

Benefits

- Understanding more about how our users use library systems, interfaces, and services helps us understand where to focus our efforts to make the improvements that most benefit them.
- Improving our data collection methods and analysis tools helps us collect more accurate, more relevant, and a higher volume of data, allowing us to form a more accurate picture of our users.
- Most current library resource discovery paradigms still require users to use the library website as a portal to discover the resources that the library offers; understanding pathways that users take to reach information resources they need that are external to the library may help us tap into an as of yet underserved group of users—those who tend not to use the library website in their workflow.
- Understanding the serendipitous resource discovery that browsing enables may help us to develop novel discovery interface components that meet needs in unique and unexpected ways.

Objective 3: Continue to work with other UNT Libraries departments to improve data quality.

Actions

- Work with the Cataloging Work Group to help find methods to batch update records to RDA.
- Investigate the possibility of using Sierra's improved access to back-end catalog data to generate better reporting mechanisms, tools, and data displays for better database maintenance.

Summary

Much of what we've undertaken as part of our RDS improvement efforts has involved discovery systems and interfaces. But bad *data* hampers discoverability just as much as poor systems and unusable interfaces—after all, systems can only implement functionality that the data supports.

Data might be considered bad for many reasons. It might be low quality because it was not entered carefully or because there were errors in the process that derived it. Or, the format in which the data is stored might not be robust enough to support desired functionality. Unfortunately, library data formats (such as MARC and AACR2) are quite old and store data in such a way that makes it difficult to parse reliably. Fortunately, new formats and standards—such as RDA and BIBFRAME—seem promising. But

we have such a large amount of data trapped in old formats that converting it all manually would be impossible; we need to develop methods and tools that will allow us to update our old data to the new formats in batch. Otherwise our data is stuck in obsolete formats that limit the improvements we can make to our systems and interfaces.

We propose that we work with those who are the experts in library data—such as members of the Cataloging and Circulation working groups—to develop, test, and implement tools that will enable us to make large-scale improvements to data quality as well as to update the format of our data in batch.

Benefits

- Improving data quality leads to better resource discoverability.
- Developing methods to batch update records to new formats (such as RDA) can help inform future efforts to transform metadata. This will help us prepare for larger format changes such as BIBFRAME that loom on the horizon. It can also help us learn how to transform metadata into novel formats to better support new discovery and user interaction models.
- Developing better tools to support database maintenance allows us to make larger data quality improvements more quickly.

Objective 4: Deploy an application to power a single-search tab on the library website that delivers a bento-box-style results display.

Actions

- Examine existing resource-discovery components and determine which should feed the results page.
- Gather technical requirements for programmatically querying, obtaining search results from, and displaying results for each desired component. (This will include gaining access to the Summon API from Serials Solutions and learning its capabilities.)
- Find a solution that best meets the technical requirements.
- Estimate total implementation time once a solution is chosen.
- Develop layout design for the results page based on existing discovery-system usage data.
- Implement the bento-box search results display and the single-search tab.
- Allow ample time for testing—including user testing—before launching the new service.

Summary

The centerpiece of our RDS interface—the global, tabbed search box that serves as a gateway to our collections—contains options for searching books & media, online articles, databases, e-journals, UNT Digital Collections, and subject guides. Although this configuration serves well users who understand our systems and collections, novice users may not know which tab they should use and may not have the time or patience to find out via trial and error.

We propose to implement a single-search, *Find All* tab on the library website search box as an addition to the existing options that would serve as the new default. Entering a query into the search box on this tab would trigger a locally-developed application that would broadcast the query to several targets—such as the catalog, Summon, the Digital Collections, Libguides, etc.—and then present the top search results from each target in its own section on one results page, using a bento-box style design.

This solution would meet the need for a single-search option within our RDS infrastructure and also address the concerns about combined searches (i.e., those that use one combined relevance-ranked results list) laid out in our original RDS report.

Benefits

- Having a single-search tool available and presenting it as the default option accommodates users who prefer a simpler resource discovery option that doesn't require them to pick a specific tab to search.
- The sectioned, "bento-box" style display is friendlier than a solid list of undifferentiated results, as there's less chance that what a user is trying to find will be drowned out by other types of results.
- It allows users to select the type of results they want *after* they've searched—they don't have to try their search multiple times in multiple search boxes if they're unsure about which tool to use.
- A locally-developed solution can be integrated seamlessly into the library website thereby improving the discovery experience.
- Implementing the single-search as one option among many (e.g., searching the catalog, searching articles, searching databases, etc.) accommodates users who may sometimes want to search everything simultaneously but may also want to select different options.

Timeline

We propose to have the four objectives outlined in the Phase Two Action Plan completed by the end of August 2015. Due to the many dependencies among objectives and tasks throughout this phase, we are

hesitant to recommend a timeline that is any more granular. As we progress, we will provide periodic status reports on our RDS blog at http://blogs.library.unt.edu/rds to ensure we keep library staff informed. We will plan to kick off another RDS reevaluation cycle in Fall 2015.

References

Bertram, A. (2011, April 7). *bertrama's sandbox: Article Discovery*. Retrieved from http://drupal.org/sandbox/bertrama/1119778

Dempsey, L. (2012, December 10). Thirteen ways of looking at libraries, discovery, and the catalog: Scale, workflow, attention. *Educause Review Online*. Retrieved from http://www.educause.edu/ero/article/thirteen-ways-looking-libraries-discovery-and-catalog-scale-workflow-attention

Jung, B. & Jones, R. (2013, April). *Sierra APIs*. Presentation at Innovative Users Group Conference, San Francisco, CA. Slides available online from http://csdirect.iii.com/documentation/presentations/iug2013_E04.ppsx

Lown, C., Sierra, T., & Boyer, J. (2013, May). How users search the library from a single search box. *College & Research Libraries, 74 (3)*. Retrieved from http://crl.acrl.org/content/74/3/227.full.pdf+html

Phillips, M., Weng, Z., Thomale, J., & Hicks, W. (2011, September 14). *Resource Discovery Systems at the UNT Libraries*. Retrieved from https://ui.library.unt.edu/project-manager/sites/default/files/resource_discovery_systems_ui_report_final.pdf

Rochkind, J. (2012, October 2). Article Search Improvement Strategy. [Web log post]. Retrieved from http://bibwild.wordpress.com/2012/10/02/article-search-improvement-strategy/

Weng, Z. (2012, May 16). *What did our Users Tell Us about How We Should Improve the Library Website Interface? And What are Our Actions in Response?* Presentation given at the University of North Texas, Denton, TX. Slides available online from http://www.library.unt.edu/events/digital-libraries/what-did-our-users-tell-us-about-how-we-should-improve-libraries-website